Trumped Innovation: Policy Uncertainty in the H-1B Strategic Factor Market*

Yang Fan† Lubomir P. Litov‡ Mishita Mehra§ Mu-Jeung Yang¶

May 2026

The rule of law in strategic factor markets is usually treated as a fixed background condition rather than a potentially changing institution. We analyze the innovation ecosystem implications of policy uncertainty from discretionary executive action that undermines predictable rules without changing them. The market for hiring global talent by US firms, governed by the H-1B visa program, offers a setting where discretionary policy changes are accurately measured and their effects on innovation are cleanly identified through the H-1B lottery. A widely underappreciated fact is that USCIS can deny visas on a discretionary basis to winners of the H-1B visa lottery, and such discretionary denials rose sharply during the first Trump administration alongside a surge in "requests for evidence." We show how this form of discretionary policy uncertainty attenuated the benefits of H-1B workers for startup innovation and VC funding, shifting the benefits of the H-1B strategic factor market away from startups and toward incumbents.

Keywords:

JEL Codes: L20; L26; L29; D21; D23

*We would like to thank the editors, seminar participants at the Leeds School of Business, University of Colorado, as well as Mukund Chari, Janet Bercovitz, Sina Sokhan and Tony Tong for useful comments.

†Department of Economics, Colby College

‡Division of Finance, Price College of Business, University of Oklahoma

§Department of Economics, University of Richmond

¶Strategy, Entrepreneurship, and Operations, Leeds School of Business, University of Colorado Boulder

1. Introduction

As the rule of law weakens in long-standing democracies such as the US (Flammer, Malhotra, and Loch, 2025a), growing executive discretion makes businesses vulnerable to policy uncertainty resulting from arbitrary administrative decisions. A key property of the rule of law is a commitment to clear "rules of the game", a commitment that can be undermined by arbitrary decisions that contradict previously understood rules. In turn, such a loss of commitment to rules introduces policy uncertainty and related unpredictability that especially hurts innovation efforts, as such investments are often firm-specific or irreversible. This effect of policy uncertainty is well understood, both in legal doctrines of "stare decisis"[1] and "reliance" on contracts (Shavell, 2004) as well as in the strategic management literatures on transaction cost economics (Williamson, 1985), and real options (Trigeorgis and Reuer, 2017). Policy uncertainty is particularly damaging amid a fierce global competition for talent, where a nation's ability to attract top innovators is a key strategic and geopolitical asset. Nowhere are these effects more prominent than in entrepreneurial ecosystems, where startups and VCs rely on predictable conditions to help world-class talent create the game-changing technologies of the future (Kerr, Nanda, and Rhodes-Kropf, 2014a; Mallaby, 2022). Indeed, the VC investor Deedy Das from Menlo Ventures recently stated: "But the top, top talent — a lot of people even I know — are reconsidering coming to America. They're like, why would I want to be in an environment where overnight I might have to fly back because an executive order is signed?" (NY Times, Dealbook Newsletter, Sept 27, 2025).

We therefore ask: What is the impact of policy uncertainty from executive discretion, especially in the area of skilled immigration, on startup innovation and VC investment? We investigate this question in the context of the H-1B visa program, the main source of high-skill global talent for the US startup ecosystem. A stylized fact that is currently under appreciated in the empirical literature on the effects of H-1B visas (Chen, Hshieh, and Zhang, 2021a; Dimmock, Huang, and Weisbenner, 2022a) is that the executive branch represented by the U.S. Citizenship and Immigration Services (USCIS) has discretion to deny H-1B visas even to firms whose petitions otherwise meet the program's statutory criteria. We document a substantial increase in such discretionary denials to H-1B visa lottery winners during the first Trump administration, and use this as a natural experiment to analyze the impact of policy uncertainty on startup innovation and VC funding. A second attractive feature of our empirical setting is that H-1B visas have been randomly allocated via a lottery in most years since 2004, thereby facilitating the estimation of causal effects of H-1B workers on startup

[1] "Stare decisis", which is often translated as "to stand by things decided", is a common law principle, according to which courts should adhere to previously set precedent and not change existing legal principles.

innovation (Chen et al., 2021a; Dimmock et al., 2022a). Together the two features of our empirical setup allow us to provide credible causal estimates on how H-1B policy uncertainty affects the innovation benefits from high-skill immigrants.

Our contribution is twofold. Theoretically, we extend strategic factor market theory by showing that the cornerstones of competitive advantage (Peteraf, 1993) rest on an implicit rule-of-law foundation. When this foundation is destabilized through discretionary administrative action, i.e. even without changes to formal rules, the cornerstones are reshaped in ways that tilt competitive advantage away from innovative entrants and toward incumbents who can shield themselves from arbitrary administrative fiat, e.g. through political lobbying. Empirically, we provide the first causal evidence that policy uncertainty differentially attenuates the causal effects of winning H-1B visas, with effects concentrated on startups rather than large incumbents. These empirical results reinforce our theoretical claim that the rule of law especially protects resource-limited innovative entrants, such as startups that may be developing game-changing innovations but lack the political connections to shield themselves from arbitrary administrative action. The rule of law is often discussed as an abstract institutional ideal, particularly in debates about democratic resilience. Our paper provides a tangible empirical illustration: when discretionary administrative action erodes predictable rules, the costs fall disproportionately on the startups that democracies most need to enable the innovations of the future and on the entrants providing competitive checks on politically connected incumbents.

2. Theory and Hypothesis Development

2.1 Strategic Factor Markets and the Rule of Law

Strategic factor market theory asks under which circumstances firms can generate sustainable rents from acquiring strategic resources (Barney, 1986; Benner and Zenger, 2016). This question has a canonical synthesis in Peteraf (1993), which identifies four conditions that jointly support sustained competitive advantage: heterogeneity, ex post limits to competition, imperfect resource mobility, and ex ante limits to competition. These four "cornerstones" organize how strategy scholars think about which resources generate rents and why.

The market for global talent in the US, as governed by the H-1B visa program, is a particularly important instance of a strategic resource market. H-1B workers are valuable (Chen, Hshieh, and Zhang, 2021b; Dimmock, Huang, and Weisbenner, 2022b), rare (quota-constrained since 2004), costly to imitate (Chen et al., 2021b; Glennon, 2024) and non-

substitutable (visa rules restrict employers from hiring foreign workers who could be replaced by comparable domestic workers). Furthermore, random allocation among oversubscribed applicants since 2008 means lottery wins are plausibly orthogonal to firm-level confounders, allowing causal rather than correlational estimation of resource effects.

But the Peteraf (1993) framework, and most empirical work on strategic resources rests on an assumption that is rarely made explicit: transactions in the strategic factor market are conducted on a foundation of the rule of law. This means that firms consider the rules of the game fixed and predictable. Contracts remain enforceable, regulatory permissions granted remain durable, and the institutional architecture of the strategic resource market is stable, Peteraf (1993). This assumption is almost never named in the literature. It is inherited from a broader tradition in economics and strategy that treats the rule of law as a background condition rather than a variable of interest. As a result, most empirical SFM work consequently focuses on market-level frictions such as asymmetric information or switching costs, rather than on institutional instability (Arts, Cassiman, and Hou, 2023; Fan, Litov, Yang, and Zenger, 2026). The few strategic factor market studies that do engage with legal institutions examine permanent shifts in stable doctrines. Flammer and Kacperczyk (2019) are a leading example: they study how a court-created doctrine (inevitable disclosure) affects firm investments in knowledge workers. Empirical studies such as this, which analyze the impact of a change in legal doctrine, are revealing precisely because they also hold the rule of law assumption fixed. The doctrine is a stable, judicially created rule whose content varies across states but whose enforcement is reliable.

But, what happens when the rule itself becomes unstable, not in the formal content of laws but in enforcement by the executive? That question is increasingly urgent and Flammer, Malhotra, and Loch (2025b) identify the erosion of rule-of-law institutions in historically democratic regimes as a first-order concern, and explicitly note that the intersection of business research with democratic resilience has been "largely left unaddressed." Our empirical setting offers an unusually rich opportunity to study this question. The first Trump administration introduced a sharp discretionary shock without changing formal rules: denial rates on lottery winners rose from 6% in FY 2015 to 24% in FY 2018 (National Foundation for American Policy, 2020), reflecting a strong increase in administrative discretionary denials of winners of the H-1B visa lottery, see Donald J. Trump (2017). This discretionary deviation from statutory laws represents a case of increased policy uncertainty, resulting from an erosion of the rule-of-law, since executive discretion results in material changes in the rules of the game without a formal change in these rules.[2]

[2]For example, H-1B lottery mechanics themselves were stable throughout the relevant period. The Department of Homeland Security (DHS)'s own description, in its 2019 rulemaking, confirms that "USCIS then may

We are interested in how changes in rule-of-law and the related increased policy uncertainty predict when the foundations of strategic factor markets change. To do this requires a theory of investment in the wake of increased policy uncertainty.

2.2 Why use Real Options to Understand Strategic Resources under Policy Uncertainty?

The traditional approach to investments, such as NPV or neoclassical investment theory (Hayashi, 1982; Jorgenson, 1963; Tobin, 1969), treats management as passive (Leiblein, Chen, and Posen, 2017): firms optimize on expected returns, with no independent channel through which policy uncertainty affects investment. But a key strategic dimension involves the tradeoff between commitment to specialized and costly-to-reverse investments (Ghemawat, 1991; Teece, 1986; Williamson, 1985) and the flexibility to wait until uncertainty resolves (McGrath, 1997; Trigeorgis and Reuer, 2017). When investments are genuinely irreversible in this sense, as is often the case with R&D and other innovation activities, uncertainty creates an option value of waiting: as information arrives, deferring preserves flexibility to respond to new information instead of committing prematurely (Dixit and Pindyck, 1994; Pindyck, 1991). For startups, this option value of waiting reduces investments in innovation efforts and the complementary assets that support innovative offerings (Teece, 1986). Additionally, it has long been recognized that VCs manage irreversible investments into startups through staged investments: Ghemawat (1991) wrote that "staged capital commitment is valuable to venture capitalists because it preserves the possibility of abandoning the venture." VCs therefore make small initial bets, such as a seed investment to fund a prototype or MVP (Ries, 2011) and abandon ventures that fail this market test.[3] For both, high-tech startups and VCs, which both deal with high degrees of uncertainty and irreversible investments, real options provides a more nuanced understanding as argued by Adner and Levinthal (2004): "to the degree that investment choices have the property of high uncertainty and irreversibility, a real options valuation provides a better characterization of the investment's true value than does a net present value calculation because of the latter's inability to account for the value of delaying commitments."

Our setting offers an unusually sharp test of the strategic implications of real options.

randomly select from the cap-subject petitions received on the final receipt date" (84 Fed. Reg. 888, 893 (Jan. 31, 2019); Department of Homeland Security, 2019); electronic pre-registration was suspended until the FY 2021 cap season (90 Fed. Reg. 60864, 60869 (Dec. 29, 2025); Department of Homeland Security, 2025). Variation in denial rates is therefore attributable to adjudicative discretion alone.

[3]While recent finance research frames staged VC funding as a sequence of "experiments" (Kerr et al., 2014a; Nanda and Rhodes-Kropf, 2017), the underlying logic of small commitments revealing informative signals predates this framing (Ghemawat, 1991).

Discretionary denial rates on H-1B lottery winners rose from 6% in FY 2015 to 24% in FY 2018 (National Foundation for American Policy, 2020), sharply increasing the risk of H-1B-dependent investment. At the same time, the Tax Cuts and Jobs Act of 2017 (TCJA) raised expected after-tax returns to productive investment, including skilled labor, so expected returns and policy uncertainty were rising simultaneously.

A traditional NPV framework predicts an increase in H-1B-dependent investment during this time: higher expected returns should increase investment, and uncertainty has no independent channel.[4] In contrast, in the real options framework (Bloom, 2009; Dixit and Pindyck, 1994; Pindyck, 1991; Sanford and Yang, 2022) policy uncertainty enters firm decisions through the option value of waiting or abandoning, not only through expected returns (Baker, Bloom, and Davis, 2016; Kim and Kung, 2017). Elevated uncertainty can therefore suppress new investment and trigger abandonment of in-progress projects even when fundamentals improve.

Two further features sharpen the predictions for this setting. Innovation projects combine market, technical, and resource-access uncertainty simultaneously, making them more fragile to exogenous shocks than routine investment (Huchzermeier and Loch, 2001; Pich, Loch, and De Meyer, 2002). And startups operate near the edge of manageable uncertainty, so additional uncertainty compounds rather than adds (Loch, Solt, and Bailey, 2008).

2.3 How Discretionary Policy Uncertainty Reshapes the Cornerstones of Competitive Advantage

Our integration of strategic factor market analysis and real options responds to recent calls for applying real-options logic to nonmarket strategy and environmental shocks (Chintakananda, McIntyre, and Tong, 2024). With this combined framework in hand, we revisit Peteraf's (1993) four cornerstones and ask how each is affected when executive discretion introduces policy uncertainty into the H-1B strategic factor market. The organizing claim is that all four cornerstones are affected, but in distinct ways, and that the real-options mechanism sharpens each prediction beyond what SFM alone would generate.

2.3.1 Heterogeneity

The first cornerstone in Peteraf's work is "resource heterogeneity, from which come Ricardian or monopoly rents". Under stable rule of law, H-1B workers generate cross-firm

[4]More precisely, neoclassical investment theory predicts that the corporate rate cut from 35% to 21% raised the after-tax marginal product of every productive input, including skilled labor (Hall and Jorgenson, 1967). Firms invest until expected after-tax MPK equals user cost of capital, so the tax-driven rise in expected MPK should raise the optimal capital stock and pull investment upward.

heterogeneity in inventive capability (Chen et al., 2021a; Dimmock et al., 2022a), engineering productivity, and scaling capability (Glennon, 2024). Policy uncertainty reshapes the source of this heterogeneity in two ways that operate at different levels. Within the firm, startups that experience discretionary denial after winning the lottery update their beliefs about the "rules of the game" and therefore about future denial risk. Because H-1B visas are valid for only three years, the relevant uncertainty is not a one-time acquisition risk but a sequence of renewal, extension, and future hiring decisions.

Traditional real options analysis in economics (Abel and Eberly, 1994; Bloom, 2009; Dixit and Pindyck, 1994) or finance (Kim and Kung, 2017; Sanford and Yang, 2022) would primarily focus on the direct impact of increase policy uncertainty on innovation investments. However, guided by the performance-focus of the strategic factor market framework, we are instead mainly interested in how policy uncertainty changes the causal effect of H-1B workers on innovation, through its impact on complementary asset (Teece, 1986) investments. To clarify this channel, consider the performance equation

$$y_{i,t} = \theta_1 \cdot H_{i,t} + \theta_2 \cdot A_{i,t} + \theta_3 \cdot (H_{i,t} \times A_{i,t}) \tag{1}$$

where, $y_{i,t}$ are patent applications, $H_{i,t}$ an indicator for having an H-1B worker and $A_{i,t}$ is a complementary asset stock. Following Milgrom and Roberts (1995), one can define complementarity of H-1B workers and assets as

$$\frac{\partial^2 y_{i,t}}{\partial H_{i,t} \partial A_{i,t}} = \theta_3 > 0 \tag{2}$$

At the same time, the policy uncertainty causing a delay in investments of complementary assets, leads to the prediction that $\frac{\partial A_{i,t}}{\partial PU_t} < 0$.[5]

Putting these two elements together leads to the interaction term:

$$\frac{\partial^2 y_{i,t}}{\partial H_{i,t} \partial PU_t} = \frac{\partial^2 y_{i,t}}{\partial H_{i,t} \partial A_{i,t}} \cdot \left(\frac{\partial A_{i,t}}{\partial PU_t} \right) = \underbrace{\theta_3 \cdot \frac{\partial A_{i,t}}{\partial PU_t}}_{<0} \tag{3}$$

In other words, the causal effect from an H-1B worker is negatively moderated in times of increased policy uncertainty because firms reduce their investment in complementary assets. At the same time, exposure to the same policy uncertainty is unlikely to be symmetric across

[5]We also note that a parallel prediction holds for demand for H-1B workers: $\frac{\partial H_{i,t}}{\partial PU_t} < 0$. This prediction especially applies to small firms and startups, as the direct costs of participating in the H-1B visa process, estimated to be roughly $20,000 are substantial irreversible costs for these firms. However, we directly exploit randomization of H-1B to firms, this potential endogeneity concern will directly be addressed by our identification strategy.

firms. Large incumbents have access to political mechanisms, such as lobbying, relationship capital, and legal resources (Hadani and Schuler, 2013; Hillman, Keim, and Schuler, 2004), which the strategy literature suggests can shield firms from policy uncertainty. On the other hand, startups typically lack such non-market resources. We cannot directly observe these mechanisms in our data, but the asymmetry implies a testable prediction: the cross-firm differential effects of policy uncertainty should be larger for startups than for large incumbents.

Hypothesis 1. H-1B discretionary policy uncertainty attenuates the positive effect of winning the H-1B lottery on startup outcomes, with the attenuation especially pronounced among smaller startups that have fewer internal resources and less capacity to absorb uncertainty.

2.3.2 Ex Post Limits to Competition

The second cornerstone of Peteraf's framework focuses on product market competition after resource acquisition, or "ex post limits to competition", which "are necessary to sustain the rents". Sustained competitive advantage from a resource can arise not only from permanent imitation barriers but from that resource generating a chain of temporary advantages, in which each generation of innovation enables the next (D'Aveni, Dagnino, and Smith, 2010; McGrath, 2013). In the H-1B setting, this chain operates through accumulated co-specialized investment around the worker: job responsibilities built around their expertise, long-horizon project commitments, firm-specific knowledge they transfer to colleagues, and client relationships anchored in their work. The lottery win grants access, then year-1 inventions enable year-2 inventions that build on the former and so on. Each invention is a temporary advantage, but the sequence can generate persistent firm-level returns. If co-specialized firm-specific assets drive persistence, policy uncertainty should reduce it because firms hesitate to invest. Under the co-specialization mechanism, policy uncertainty breaks the chain through real-options logic: facing uncertain retention beyond year 3 (the time horizon of an H-1B visa), firms defer co-investments and abandon in-progress innovation rather than commit irreversibly (Dixit and Pindyck, 1994; Ghemawat, 1991; Pindyck, 1991). Fewer year-1 patents mean a weaker foundation for year-2 patents and so on. This mechanism is testable against an alternative. If the worker's contributions are non-rival ideas (Arrow, 1962; Fan et al., 2026) the firm can keep building on as long as the worker is present, persistence should be largely unaffected within the three-year visa window.

Hypothesis 2. High discretionary H-1B denial rates reduce the persistence of past H-1B lottery wins on current startup innovation, because the firm cannot sustain the co-investment sequence around H-1B workers long enough for the innovation chain to compound. Absence

of such attenuation would suggest that persistence operates through non-rival contributions of the worker rather than through accumulating firm-specific co-investment.

2.3.3 Imperfect Resource Mobility

Peteraf's third cornerstone focuses on resource mobility, since "Imperfect resource mobility ensures that the rents are bound to the firm and shared by it". In the H-1B setting, visa rules tie workers to sponsoring firms, supporting this cornerstone institutionally. During the first Trump administration, this institutional protection was further reinforced. U.S. Citizenship and Immigration Services (2017) instructed adjudicators to re-evaluate H-1B extension and transfer petitions from scratch rather than deferring to prior approvals. Lawyers immediately recognized this as raising the bar for previously routine extensions and transfers, and elevated denial rates followed (National Foundation for American Policy, 2018, 2020). For workers already at the firm, switching employers became materially harder. This should make a discretionary denial for continuing H-1B workers more costly: Transfer petitions face the same scrutiny as new petitions and denied workers are more likely to leave the country entirely, while the firm cannot easily replace them with another H-1B worker. However, this intuition misses the dynamic perspective and what firms do before a denial occurs. Co-specialized investment around an H-1B worker, such as embedding them in firm-specific projects, building specialized roles, transferring firm-specific knowledge are irreversible and only pay off if the worker remains. Under elevated continuation denial risk, the option value of deferring such commitment rises (Dixit and Pindyck, 1994; Ghemawat, 1991; Pindyck, 1991). The structural problem is familiar from transaction cost economics (Folta, 1998; Williamson, 1985): firms underinvest in transaction-specific assets when those investments are vulnerable to loss, whether the threat is opportunistic holdup by a counterparty or, as here, administrative removal of the counterparty entirely. Firms will therefore withhold co-specialized investments to H-1B workers, assigning workers to generic roles, not building specialized teams around them and not transferring firm-specific knowledge.

The withholding response should be most pronounced at small firms. Their productive scale is small relative to the cost of failed commitment, so the threshold for committing to co-specialized investment is lower. Small firms switch to withholding at denial probabilities where larger firms still commit (Appendix A formalizes this). This discussion highlights the value of using real options logic as our theoretical framework to analyze resource mobility, rather than just using a static approach. Importantly, the predicted positive interaction does not reflect uncertainty being benign. The firm's expected payoff is lower under high uncertainty because it operates at the lower productive value of generic deployment. Only the incremental response to specific denial events appears smaller.

Hypothesis 3. The incremental damage of a continuation denial on startup innovation is smaller during periods of high discretionary H-1B denial rates than during periods of low discretionary denial rates, because under elevated uncertainty, firms will tend to withhold the co-specialized investment that denials would otherwise destroy. The effect should be stronger at small firms than at large incumbents.

2.3.4 Ex Ante Limits to Competition

Peteraf's fourth cornerstone states that "Ex ante limits to competition prevent costs from offsetting the rents", i.e. firm does not pay away future rents at resource acquisition. The H-1B lottery serves the role of resource acquisition: random allocation prevents auction-style bidding among firms. The lottery itself was structurally stable through our sample period (84 Fed. Reg. 888, January 2019; Department of Homeland Security, 2019), so the cornerstone's institutional protection is unchanged.

But for high-growth technology startups as those covered in our data, the relevant ex ante competition is between VCs backing startups, not between firms bidding for the worker. Whether the startup captures rents from H-1B-enabled innovation depends on VCs continuing to fund the firm at terms that allow the worker to be deployed profitably. In this context, VC funding is itself a sequence of real options. Each round is a staged, conditional commitment with the option to continue or abandon at the next stage (Ghemawat, 1991; Kerr, Nanda, and Rhodes-Kropf, 2014b; Nanda and Rhodes-Kropf, 2017). VCs price not just expected returns but also the value of preserving these continuation options.

Elevated H-1B policy uncertainty raises that option value through two distinct channels at startups exposed to discretionary denials. First, when an H-1B lottery winner is denied at the petition stage, the firm loses the worker before any productive deployment, and the immediate growth path the startup pitched to its VCs is disrupted. VCs revise downward both the expected return and their belief about the startup's ability to execute under the prevailing policy regime, which raises the option value of deferring further funding. Second, even past winners, i.e. startups that successfully obtained H-1B workers within the past three years, will become riskier bets. The mechanisms in Hypotheses 1–3 (attenuated benefit, broken persistence, withholding of co-investment) imply that the future productive value of past wins is more uncertain than it appeared at the time of acquisition. VCs facing this elevated continuation risk preserve their abandonment options by reducing funding. In both cases, real-options logic explains the response: elevated uncertainty about the productive value of H-1B workers raises the option value of deferring or abandoning continuation, and VCs preserve this option by reducing investment.

Hypothesis 4. During periods of high discretionary H-1B denial rates, VC funding

falls (a) for startups whose H-1B petitions are denied after winning the lottery, and (b) for startups that won the lottery within the past three years.

3. Empirical Methodology

Our empirical strategy exploits two key sets of features. On the one hand, the institutional setting of the H-1B visa program facilitates the credible identification of causal effects. H-1B visas have been allocated through a lottery in most years since 2004, which provides a plausibly exogenous source of variation in firms' access to high-skill foreign labor (Chen et al., 2021a; Dimmock et al., 2022a). The process of applying for an H-1B visa is well-known to be time intensive and costly, with estimates putting the pure financial cost at roughly $ 20,000 per application. A less appreciated fact is that winning the H-1B visa lottery still does not entitle an employer to hire an H-1B worker, since the executive branch as represented by the US Citizenship and Immigration Service (USCIS), and importantly, appointed by the current presidential administration, has discretion to deny H-1B visas to winners of the H-1B lottery. Importantly, the denial of new hires likely implies only a small direct search cost, but it is a strong signal about the unpredictability of immigration policy, which in turn disproportionately increases the perceived uncertainty of dealing with the Trump administration.

On the other hand, the analysis of H-1B changes during the first Trump administration as a natural experiment offers several advantages. First, Trump's election victory was surprising to most observers, which rules out prior preparation of entrepreneurs or venture capitalists. Second, policy changes were not in anticipation of a recession, which is the typical endogeneity problem in the analysis of other government policy changes, such as tax reform (Romer and Romer, 2010). And indeed, there was no recession until the unanticipated and short-lived COVID-19 pandemic recession in the last year of Trump's presidency in 2020. This matters, since in many settings in economics and finance, the effects of increased uncertainty are often confounded with the effects of bad news or low expected returns (Baker et al., 2016; Bloom, 2009; Kim and Kung, 2017). In contrast, as we discussed in section 2.2, the TCJA qualifies as what Romer and Romer (2010) call an "exogenous" tax shock, that stimulated investment and if anything will tend to bias results against finding an effect of policy uncertainty on the H-1B strategic factor market.

Our main specification takes the form:

$$y_{i,t} = \beta_1 \cdot T_{i,t} + \beta_2 \cdot (\delta_t \times T_{i,t}) + D_t + D_i + \epsilon_{i,t} \tag{4}$$

where $y_{i,t}$ is an outcome for firm i in year t, $T_{i,t}$ is a treatment indicator equal to one if the firm in that year is an H-1B lottery winner, δ_t is the annual aggregate fraction of lottery winners that are discretionarily denied an H-1B visa, D_t is a full set of time fixed effects, D_i is a full set of firm fixed effects, and $\epsilon_{i,t}$ is an error term. The baseline effect $\beta_3 \cdot \delta_t$ is being absorbed by the time fixed effects D_t. Standard errors are clustered at the firm level.

This specification maps closely to the economic mechanism motivating the paper and shown in equation (1). The coefficient β_1 captures the average treatment effect of lottery success on firm outcomes. The time fixed effects D_t control for macroeconomic policies such as the TCJA. Additionally, time fixed effects absorb the direct effect of policy uncertainty on firm outcomes, such as patent applications or VC funding. This is typically the effect of interest in economics and finance studies, such as Baker et al. (2016); Kim and Kung (2017). In contrast, our main coefficient of interest is β_2, which captures treatment effect heterogeneity in the value of winning the H-1B lottery as a function of aggregate discretionary denials. A negative estimate of β_2 implies that the innovation and financing benefits of lottery success are attenuated when the surrounding policy environment becomes less predictable. In other words, even if firms win the lottery, the value of that win declines if executive discretion makes ultimate approval less credible.

We examine this relationship across two broad classes of dependent variables. The first consists of innovation outcomes, namely patent applications and patent grants. Patent applications capture contemporaneous innovative activity, while patent grants are measured with a three-year lag in order to better capture successful or higher-value innovation that survives the patent examination process. Because these variables are count outcomes, we estimate the relevant specifications using Poisson panel regressions with firm and year fixed effects. The second class consists of financing outcomes, including total funding, VC funding, and Series A funding. These outcomes are estimated using linear regressions with firm and year fixed effects. This empirical design allows us to evaluate whether H-1B policy uncertainty affects both the production of innovation and the financing environment surrounding startups.

A key question is whether the baseline interaction effect reflects the broader uncertainty created by discretionary denials or simply the direct consequences of realized denials at firms that win the lottery. To address this distinction, we estimate within-winner specifications that replace the baseline treatment indicator with alternative measures of realized denial exposure. These include an indicator for firms with any denials, an indicator for firms with all lottery-winning applications denied, and a continuous firm-level denial rate. These regressions allow us to assess whether firms respond primarily to their own realized denial outcomes or instead to the broader risk that a lottery win may no longer reliably translate into an H-1B worker.

To investigate Hypothesis 3, we also distinguish between denials involving new H-1B workers and denials involving continuing workers whose visas must be renewed after three years. Let $C_{i,t}$ denote an indicator equal to one if firm i experiences a denial of a continuing H-1B application in year t. We then estimate specifications of the form

$$y_{i,t} = b_1 \cdot C_{i,t} + b_2 \cdot (C_{i,t} \times \delta_t) + D_t + D_i + \epsilon_{i,t} \tag{5}$$

where again as previously, the baseline effect $b_3 \cdot \delta_t$ is being absorbed by the time fixed effects D_t. From our discussion of Hypothesis 3, one would expect that $b_2 > 0$ because losing an H-1B worker is less damaging to a firm that has withheld complementary asset investments under policy uncertainty. However, the decision to file for H-1B continuation is itself endogenous to firm behavior, creating a selection problem. A firm that has withheld co-specialized investment in its H-1B worker has less to lose from letting the worker go, so it is less likely to file for continuation under high uncertainty. The applicants we observe under high uncertainty are therefore disproportionately firms that actually committed to co-specialized investment. This biases b_2 downwards, working against the H3 prediction.

To further examine heterogeneity in exposure to H-1B policy uncertainty, we construct a set of firm-level measures based on prior participation in the H-1B system. These include indicators for whether the firm was a lottery winner in the previous year, whether it was a winner two years earlier, and whether it has ever won an H-1B lottery. These specifications are motivated by the idea that firms with prior success in the H-1B market are more likely to rely on high-skill foreign labor and therefore more likely to be exposed when discretionary denial risk rises. If so, then the interaction between these past-winner indicators and the aggregate denial rate should be negative for innovation and financing outcomes.

We also analyze more persistent measures of participation in the H-1B market. In one set of specifications, we classify firms according to whether they always apply for H-1B visas or never apply. In another, we refine this approach by measuring application intensity prior to 2017, before the sharp increase in discretionary denials. Specifically, we classify firms as always apply, single apply, or never apply based on their pre-2017 H-1B application behavior. These pre-period measures are useful because they capture baseline exposure to the H-1B system before the policy environment changes. Interacting these indicators with the aggregate denial rate allows us to test whether firms with stronger prior reliance on H-1B hiring experience larger declines in innovation and financing when discretionary denials rise.

Across all specifications, the inclusion of firm fixed effects ensures that identification does not come from permanent differences across firms in innovative capacity, funding access, or labor demand. Instead, the estimates are identified from within-firm changes over time

combined with variation in aggregate discretionary denials and differential exposure to the H-1B system. Year fixed effects absorb common shocks affecting all firms, including changes in macroeconomic conditions, venture capital markets, and the overall startup environment. In this sense, our empirical approach is designed to isolate whether startups that are more exposed to the H-1B channel are disproportionately affected when executive discretion makes the immigration process less predictable.

Taken together, this empirical methodology allows us to test not only whether access to H-1B workers matters for startup outcomes, but also whether the value of that access depends on the credibility of the surrounding policy regime. This is the central empirical implication of the paper: if executive discretion undermines the predictability of the H-1B process, then the causal benefits of winning the lottery should be weaker precisely when discretionary denials become more prevalent.

4. Data

Our sample data is combined from four data sources: startup-level information from Crunchbase, H-1B petition outcomes from U.S. Citizenship and Immigration Services (USCIS), patent records from the USPTO/PatentsView, and Labor Condition Application (LCA) filings from the U.S. Department of Labor. Together, these sources allow us to construct a firm-year panel that links startup characteristics, skilled-immigration demand and outcomes, innovation activity, and financing outcomes.

We begin with startup firms from the Crunchbase organizations database. Crunchbase provides information on firm identity and background characteristics. We restrict attention to U.S.-based firms and define startups as firms with non-missing founding dates in 2000 or later. This restriction yields a sample aligned with the entrepreneurial firms most likely to rely on external financing and high-skill labor during the period we study.

To measure firms' exposure to high-skill immigration, we use employer-level H-1B petition data from the USCIS H-1B Employer Hub spanning 2009–2021. These data report the number of approved and denied petitions for both new and continuing H-1B workers, along with employer names, state and ZIP code information, the last four digits of tax identifiers, and self-reported NAICS codes. We exclude non-profit and cap-exempt institutions such as universities, colleges, and religious organizations, since these employers are not subject to the H-1B lottery and therefore do not face the same source of uncertainty studied in this paper.

Because the USCIS data do not contain a stable firm identifier and employer names are self-reported, a substantial amount of cleaning is required before matching. We stan-

dardize employer names and aggregate petitions submitted by the same employer within each year where appropriate. We then match H-1B employers to Crunchbase firms using probabilistic record linkage based on firm-name similarity, supplemented by state and ZIP code verification. Our final matched sample includes exact name matches with consistent state information as well as manually verified high-similarity matches that also align on location. This procedure produces 48,448 Crunchbase–H-1B firm matches, of which 25,570 are confirmed startups based on the post-2000 founding-date restriction.

We next merge patenting outcomes to the Crunchbase-H-1B sample. For patent grants, we use the USPTO PatentsView bulk files, which provide patent application dates, granted patent identifiers, disambiguated assignee names, assignee locations, and patent classification information. We harmonize assignee names using the same standardization approach applied to Crunchbase and USCIS employers, and then match patent assignees to the already harmonized Crunchbase-H-1B firm sample. Since patent assignees can appear in multiple locations, we retain unique assignee-state combinations during the matching process. For patent applications, we use the USPTO bulk files on pre-grant published patent applications and apply an analogous matching procedure. The resulting Crunchbase-patent-H-1B matched sample includes 10,238 unique Crunchbase firm identifiers.

A central challenge in the H-1B setting is that the USCIS petition data only record adjudicated petitions. They do not directly identify firms that attempted to hire through the H-1B cap season but had no petition selected in the lottery. To recover this margin, we use employer-level LCA filings for fiscal years 2013-2019. The LCA is a required upstream filing for employers seeking to hire H-1B workers and therefore provides a measure of intended labor demand rather than realized visa allocation. We restrict the LCA data to certified H-1B cases, remove likely cap-exempt and non-profit employers using the same name-based screens applied to the USCIS sample, standardize employer names, and link LCA employers to the matched Crunchbase-USCIS-patent firm sample using probabilistic record linkage based on names and location.

We use the LCA filings to proxy for firm demand during the H-1B lottery season. Following the institutional timing of the H-1B cap process, we identify cap-season filings as those submitted between January and April with intended employment start dates between August and October. We then construct a firm-year measure of new H-1B applications from these cap-season LCAs. Prior to fiscal year 2017, when the LCA extracts do not consistently report new-employment counts, we proxy new applications using total positions requested on cap-season filings. Beginning in fiscal year 2017, the LCA disclosures separately report new-employment positions, and we therefore use those counts directly. Combining this LCA-based application proxy with USCIS adjudication outcomes allows us to distinguish lottery

winners from lottery losers. We define a firm-year as a *lottery winner* if it has any adjudicated new petitions, that is, at least one new H-1B approval or denial. Importantly, both approvals and denials occur only after a petition is selected in the lottery, so we treat firms with either outcome as lottery winners. We define *lottery losers* as firm-years with positive cap-season applications in the LCA data but zero new approvals and zero new denials in the USCIS data. Firm-years with missing USCIS adjudication information are left undefined. Our main measure, *Fraction Denied*, is defined as the annual ratio of denied to total adjudicated (approved plus denied) new H-1B petitions among lottery-selected applications. Our main analysis focuses on the period 2013–2019, which reflects the overlap between LCA data availability and the pre- and post-policy-change periods.

This distinction is important because it allows us to separate firms that sought to hire through the H-1B system from those that successfully entered the adjudication stage. In turn, it lets us exploit both the quasi-random allocation generated by the lottery and the subsequent uncertainty introduced by discretionary denials among lottery winners. This is the key institutional feature underlying our empirical design.

Our main outcome variables fall into two categories. The first category captures innovation. We measure contemporaneous innovation using annual patent applications and more successful or higher-value innovation using patent grants observed with a three-year lag. The second category captures financing outcomes, including total funding, VC funding, and Series A funding, all drawn from Crunchbase. In the regression analysis, the funding variables are used in log form, while the innovation outcomes are modeled as counts.

Table 1 reports summary statistics for the principal variables used in the analysis. In the full firm-year panel, firms average just under one patent application per year and about 0.7 patent grants measured with a three-year lag. Funding outcomes are highly skewed, with substantial mass at zero and a small number of very large financing events, motivating the use of logarithmic transformations in the main specifications. On the immigration side, firms receive on average more approvals than denials, but denial exposure rises sharply over the sample period, particularly during the later years of the first Trump administration. This time-series increase in discretionary denials is the core source of variation we use to study how policy uncertainty attenuates the value of H-1B lottery success.

Table 1. Summary Statistics

Variable	Obs.	Mean	Std. Dev.	Min	Max
Panel A: Outcome Variables					
Patent Applications	77,686	0.938	3.853	0	356
Patent Grants (3-year lag)	77,686	0.713	2.792	0	159
Total Funding (millions USD)	53,277	3.970	25.700	0	1,500
VC Funding (millions USD)	53,277	1.917	16.700	0	1,500
Series A Funding (millions USD)	53,277	0.299	2.753	0	100
Log Total Funding	53,277	3.215	6.288	0	18.293
Log VC Funding	53,277	1.138	4.152	0	17.728
Log Series A Funding	53,277	0.380	2.404	0	15.956
Panel B: H-1B Approval and Denial Variables					
New H-1B Approvals	28,980	0.528	3.993	0	322
Continuing H-1B Approvals	24,108	1.177	8.778	0	644
New H-1B Denials	28,980	0.040	0.745	0	102
Continuing H-1B Denials	24,108	0.041	0.405	0	31
New Denial Rate	5,597	0.066	0.216	0	1
Continuing Denial Rate	6,595	0.035	0.151	0	1
Total Denial Rate	9,110	0.050	0.179	0	1

Notes: This table reports summary statistics for the main innovation, funding, and H-1B approval and denial variables used in the analysis. Patent Grants (3-year lag) measures patents that are eventually granted with a three-year lag. Funding variables are measured in millions of U.S. dollars. The log funding variables are winsorized versions of the corresponding funding measures used in the main regressions. Denial rate variables are defined as the fraction of applications denied within the relevant category.

A final caveat is that LCA filings should be interpreted as proxies for intended H-1B demand rather than one-to-one measures of realized petitions. A certified LCA does not guarantee that the employer ultimately submits a petition for every listed position, nor that any submitted petition is selected in the lottery or approved by USCIS. For our purposes, however, this feature is useful: the LCA captures firms' ex ante demand for high-skill foreign labor, while USCIS outcomes capture the realized allocation and adjudication process. The combination of the two allows us to study not only whether firms seek H-1B talent, but also whether rising discretionary denials weaken the expected benefits of obtaining access to that talent.

5. Results

In the following sections we systematically follow the structure of our hypothesis development to document the empirical results.

5.1 H-1B Wins Under Policy Uncertainty: Heterogeneous Effects Across Firms (Hypothesis 1)

Our baseline results in Table 2 show that the positive effect of winning the H-1B lottery on startup innovation is substantially weakened during periods of high discretionary policy uncertainty. Startups that are lottery winners increase their innovation efforts as measured by more patent applications as well as more innovation outputs, as measured by more granted patents over the next 3 years. Specifically, the interaction between Lottery Winner and Fraction Denied is negative in every specification and statistically significant across columns (1) to (4).

The magnitudes are economically meaningful. In column (1), a one-percentage-point increase in the aggregate denial rate reduces the treatment effect of winning the lottery on patent applications by roughly 3.4 percent. In column (2), the corresponding reduction for granted patents is about 3.9 percent. These estimates indicate that the gains from lottery success are highly sensitive to the broader policy environment. The aggregate effect over the 2016 to 2018 period is substantial. Aggregate denial rates rose from 6.7 percent in 2016 to 10.2 percent in 2017 and then to 20 percent in 2018, so by 2018 roughly one in five lottery winners was discretionarily denied an H-1B visa. Applying the baseline interaction estimates, this 13-percentage-point increase implies approximately a 35 to 40 percent reduction in the innovation gains associated with lottery success. Much of the benefit of winning the lottery is undone when firms expect that approvals may still be reversed or denied for discretionary reasons.

The results are consistent with the real-options mechanism developed in Section 2 and shown in equation (3): Even conditional on winning the lottery, firms cannot rely on actually deploying the worker they planned to hire, which raises the option value of deferring complementary investments rather than committing to them. Startups respond by scaling back the forward-looking, irreversible investments in projects, teams, and capabilities which in turn reduces the performance benefits from winning H-1B visas.

Both our theoretical analysis and our empirical interpretation focused on market-specific policy uncertainty, due to increased discretionary denials of H-1B visa applications. However, an important concern with this interpretation may be that policy uncertainty was elevated in a more general way during the first Trump administration. Indeed, observers such as

Table 2. H-1B Policy Exposure, Economic Policy Uncertainty, and Innovation

	Baseline		Control for EPU		Never Apply		
	(1) Patent Apps	(2) Patent Grants	(3) Patent Apps	(4) Patent Apps	(5) Patent Apps	(6) Patent Apps	(7) Patent Apps
Lottery Winner	0.468***	0.426***	0.845***	0.418***			
	(0.117)	(0.118)	(0.243)	(0.149)			
Lottery Winner × Fraction	−3.398***	−3.911***	−2.329**	−3.496***			
Denied	(0.926)	(0.863)	(1.097)	(0.986)			
Lottery Winner × EPU			−0.004*				
			(0.002)				
Lottery Winner × Migration				0.0002			
Policy Uncertainty				(0.0004)			
Never Apply × Fraction					3.379***	2.255***	3.132***
Denied					(0.615)	(0.603)	(0.617)
Never Apply × EPU						0.003***	
						(0.001)	
Never Apply × Migration							0.0005**
Policy Uncertainty							(0.0002)
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	8655	8896	8655	8655	61131	61131	61131

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: Lottery Winner is an indicator equal to 1 if the firm annually has either a new H-1B approval or denial. Never Apply is an indicator equal to 1 if the firm did not apply for an H-1B visa before 2017. Fraction Denied is an annual aggregate measure equal to the number of lottery winners denied an H-1B visa divided by the total number of lottery winners. EPU is the annual average of the news-based economic policy uncertainty index. Migration Policy Uncertainty is the annual average of the news-based migration policy uncertainty index. Columns (1)–(2) report the baseline innovation specifications. Columns (3)–(7) report patent-application specifications with alternative uncertainty interactions. All specifications use Poisson panel regressions with firm and year fixed effects. Standard errors are clustered at the firm level.

Nicholas Bloom noted that Trump "changes his mind often and changes his advisors even faster. That makes him the most unpredictable president in recent times". To separate the effects of H-1B policy uncertainty from more general forms of policy uncertainty, we control for the interaction of H-1B lottery wins with economic policy uncertainty (EPU, see Baker et al. (2016)) and more general migration policy uncertainty[6]. As columns (3) and (4) show, controlling for general economic policy uncertainty or more general migration policy uncertainty leaves most the quantitative magnitude of our main estimate intact.

An additional way to rule out the possibility that our empirical results are picking up an environmental shock that is unrelated to the H-1B strategic factor market is to focus on firms that never applied to the H-1B visa program. These "Never Apply" firms should not exhibit any negative effects in times of elevated policy uncertainty due to higher discretionary denials. On the contrary, such "Never Apply" firms might actually benefit, as their H-1B dependent competitors are reducing complementary asset investments, thereby reducing the intensity of competition. Columns (5) to (7) show that this is indeed the case. What is more, these effects are not driven by either economic policy uncertainty as controlled for in column (6) nor general migration policy uncertainty as controlled for in column (7).

An important additional component of Hypothesis 1 is the heterogeneity of H-1B lottery treatment effects across firm types. As discussed in the introduction, the rule of law functions especially as protection for resource-limited innovative entrants. Larger incumbents can shield themselves from arbitrary administrative action through political resources that startups typically lack. If this protective function matters, policy uncertainty from discretionary denials should hurt smaller, less politically connected firms disproportionately, while leaving larger incumbents relatively unaffected. Table 3 shows that this is indeed the case in the data. Small firms exhibit 2-3x more negative interaction effects from H-1B policy uncertainty than large firms.

[6]General migration policy uncertainty in the US during the first Trump administration was more focused on illegal immigration through the Southern Border, and was therefore not necessarily related to high-skill immigration. For more methodological details on how general US migration policy uncertainty is measured, see: $https://www.policyuncertainty.com/immigration\_fear.html$, using the methodology of Baker et al. (2016).

Table 3. Heterogeneous Effects of H-1B Lottery and Denial Rates by Firm Size

	Panel A: Large Firms		***Panel B: Small Firms***	
	Patent Applications	Patent Grants	Patent Applications	Patent Grants
Lottery Winner	0.273	0.065	0.472***	0.119
	(0.295)	(0.285)	(0.206)	(0.087)
Lottery Winner × Fraction	−1.862	−1.338	−3.334***	−4.096***
Denied	(2.367)	(2.215)	(0.689)	(0.759)
Observations	1452	1471	7203	7425
Firm fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: Panel A restricts the sample to larger firms that have more than 500 employees. Panel B restricts the sample to smaller firms that have 499 or fewer employees. Lottery Winner is an indicator equal to 1 if the firm annually has either a new H-1B approval or denial. Fraction Denied is an annual aggregate measure equal to the number of lottery winners denied an H-1B visa divided by the total number of lottery winners. Dependent variables measure firm-level innovation outcomes, including patent applications and patent grants. All specifications use Poisson panel regressions with firm and year fixed effects. Standard errors are clustered at the firm level.

5.2 Persistence of Past Lottery Wins Under Policy Uncertainty (Hypothesis 2)

Whereas Hypothesis 1 examined the contemporary effect of winning the lottery, Hypothesis 2 examines persistence: whether past lottery wins continue to generate innovation in subsequent years. The theoretical claim from Section 2.3.2 is that this persistence operates through a chain of co-specialized investments — year-1 outputs enabling year-2 outputs — that compounds the worker's contribution over time. If policy uncertainty causes firms to defer co-investments, this chain should break. Past lottery wins should then have weaker effects on current innovation in periods of high discretionary uncertainty.

Table 4 shows that this is indeed the case[7]: past lottery winners generate fewer patent

[7]Recall that lottery assignment is independent across years, so conditioning on lagged dependent variables is unnecessary for identification. It would also be inappropriate: past innovation outcomes are themselves affected by past lottery wins and policy uncertainty, making them a "bad control" in the sense of Angrist and Pischke (2009, Chapter 3).

applications and fewer granted patents, exactly as predicted by the withholding mechanism developed in Section 2.3.2.

From a static demand-and-supply perspective, this pattern is initially surprising. Increased discretionary denials reduce the supply of new H-1B workers, which should increase the marginal value of H-1B workers who have received approval within the past three years. If, for example, the pool of H-1B workers is helpful in a similar set of innovations, one would expect less competition from new H-1B workers to result in more innovations by incumbent H-1B workers.

Another notable pattern in panels C and D of Table 4 is that the negative interaction effects become stronger past the 2-year mark. This is consistent with a tightening institutional constraint over the visa's three-year term. As a worker approaches the continuation decision, the firm's substitution options narrow: replacing a worker through a new lottery becomes harder under high uncertainty, and securing continuation also becomes harder. The marginal cost of policy uncertainty therefore rises as the worker approaches the end of the initial visa term.

This temporal pattern points us toward the continuation decision itself as the empirical focus of the next analysis. Hypothesis 3 examines what happens when firms actually face a continuation denial, and how the effect of such denials on innovation outcomes varies with discretionary policy uncertainty.

Table 4. Lagged H-1B Lottery Winners and Exposure to Denial Rates

	Patent Applications	Patent Grants
Panel A: One-Year Lagged Lottery Winner		
L1 Lottery Winner × Fraction Denied	−2.871**	−3.233***
	(1.205)	(0.936)
Observations	7053	7152
Panel B: Two-Year Lagged Lottery Winner		
L2 Lottery Winner × Fraction Denied	−3.256***	−4.172***
	(1.227)	(0.935)
Observations	5205	5267
Panel C: Three-Year Lagged Lottery Winner		
L3 Lottery Winner × Fraction Denied	−7.063***	−5.197***
	(1.604)	(0.993)
Observations	3621	3571
Panel D: Four-Year Lagged Lottery Winner		
L4 Lottery Winner × Fraction Denied	−6.134***	−5.138***
	(2.075)	(1.159)
Observations	2220	2232
Firm fixed effects	Yes	Yes
Year fixed effects	Yes	Yes

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: This table reports specifications interacting lagged H-1B lottery-winner indicators with the aggregate H-1B denial rate. L1, L2, L3, and L4 Lottery Winner indicate whether the firm was a lottery winner one, two, three, or four years prior, respectively. Fraction Denied is an annual aggregate measure equal to the number of lottery winners denied an H-1B visa divided by the total number of lottery winners. Dependent variables measure firm-level innovation outcomes, including patent applications and patent grants. All specifications use Poisson panel regressions with firm and year fixed effects. Standard errors are clustered at the firm level.

5.3 The Effect of Continuation Denials Under Policy Uncertainty (Hypothesis 3)

The previous section showed that the marginal cost of policy uncertainty grows as a worker approaches the end of the initial visa term, pointing toward the continuation decision as the moment when firm options narrow most. Hypothesis 3 examines this turning point

directly and provides discriminating tests between two mechanisms.

The intuitive mobility-friction prediction is that continuation denials hurt more under high uncertainty, since replacement becomes harder. Section 2.3.3 argues the opposite. If firms anticipate elevated denial risk by withholding co-specialized investment, then the marginal damage of an actual denial is smaller under high uncertainty, because there is less firm-specific investment at stake. The two mechanisms predict opposite signs on the lottery $\times$ uncertainty interaction.

The mechanisms also predict different patterns when comparing denied firms to firms that never won the lottery. A pure mobility-friction story implies that denied firms should perform worse than never-won firms, because they have lost a worker who was being deployed productively. The withholding mechanism predicts that denied firms should look comparable to never-won firms, because the firm had already reduced firm-specific investment in the worker. We test both predictions in the analyses that follow.

Table 5. Continuing H-1B Denials and Innovation Outcomes

	Patent Applications	Patent Grants
	(1)	(2)
Panel A: Small Firms		
Continuing Denial	-1.066**	-1.388***
	(0.451)	(0.475)
Continuing Denial × Fraction Denied	5.619*	7.769***
	(3.361)	(2.977)
Observations	209	224
Panel B: Large Firms		
Continuing Denial	0.717	0.115
	(0.531)	(0.307)
Continuing Denial × Fraction Denied	-1.186	-0.461
	(4.425)	(2.532)
Observations	228	234

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: Firm fixed effects Poisson regressions. Dependent variables are patent applications and patent grants (3-year lag). Continuing Denial equals one for firm-years with at least one continuing H-1B denial. Fraction Denied is the aggregate annual denial rate. All specifications include firm and year fixed effects. Robust standard errors are reported in parentheses.

Panel A of Table 5 shows that small startups experience less marginal damage from continuation denials when H-1B policy uncertainty is high. This is consistent with the withholding mechanism developed in Section 2.3.3: anticipating elevated denial risk, these firms reduce co-specialized investment in their H-1B workers, so denials destroy less firm-specific value when they actually occur. The pattern differs for large incumbents (Panel B). The coefficients on the (continuation denial) $\times$ uncertainty interaction are qualitatively consistent with denials being more damaging under high uncertainty, the prediction of a pure mobility-friction story, although they are not statistically significant. While the lack of significance limits what can be inferred from Panel B alone, the qualitative pattern across the two panels is consistent with the firm-size heterogeneity prediction from Section 2.3.3 and Appendix A: small firms switch to withholding at lower denial thresholds, while larger firms remain in a commit regime where denials are more costly when they occur.

5.4 VC Funding Effects of H-1B Wins Under Policy Uncertainty (Hypothesis 4)

The first three hypotheses examined how discretionary uncertainty reshapes Peteraf's first three cornerstones: heterogeneity, ex post limits, and resource mobility. Hypothesis 4 turns to the fourth cornerstone, ex ante limits to competition. Although the H-1B lottery itself remains a structurally stable allocation mechanism, Section 2.3.4 argued that the cornerstone has a second location for VC-backed startups: the relevant ex ante competition is between VCs at the funding stage, not between firms bidding for the worker. VC funding is itself a sequence of real options, where each round preserves the option to continue funding or to abandon at the next stage. Elevated H-1B policy uncertainty should raise the option value of deferral for VCs evaluating exposed startups, since the future productive value of H-1B-enabled innovation has become more uncertain. Hypothesis 4 predicts that VC funding falls during periods of high discretionary uncertainty, even for startups that won the lottery within the past three years.

Table 6 confirms this prediction. Winning the H-1B visa lottery during times of high H-1B policy uncertainty provides a smaller funding boost than under low uncertainty. This finding reflects a smaller value gain for funded startups from winning the H-1B lottery and is consistent with the withholding of complementary investments to H-1B workers documented in prior sections. Importantly, these funding effects are persistent across time periods. This persistence supports the real-options and complementary-investment channel as the dominant mechanism, ruling out the static-substitution alternative under which constrained H-1B supply should have increased the value of incumbent lottery winners.

Table 6. Lagged H-1B Lottery Winners and Exposure to Denial Rates: Funding Outcomes

	Total Funding	VC Funding	Series A Funding
Panel A: Current Lottery Winner			
Lottery Winner ×	−26.409***	−15.363***	−5.694***
Fraction Denied	(2.879)	(2.266)	(1.153)
Observations	9884	9884	9884
Panel B: One-Year Lagged Lottery Winner			
L1 Lottery Winner ×	−21.096***	−12.661***	−3.259***
Fraction Denied	(2.864)	(2.154)	(0.901)
Observations	8610	8610	8610
Panel C: Two-Year Lagged Lottery Winner			
L2 Lottery Winner ×	−19.714***	−8.628***	−4.667***
Fraction Denied	(2.992)	(2.138)	(1.075)
Observations	6847	6847	6847
Panel D: Three-Year Lagged Lottery Winner			
L3 Lottery Winner ×	−14.377***	−7.783***	−2.897**
Fraction Denied	(3.444)	(2.450)	(1.341)
Observations	5238	5238	5238
Firm fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: This table reports specifications interacting current and lagged H-1B lottery-winner indicators with the aggregate H-1B denial rate. L1, L2, and L3 Lottery Winner indicate whether the firm was a lottery winner one, two, or three years prior, respectively. Fraction Denied is an annual aggregate measure equal to the number of lottery winners denied an H-1B visa divided by the total number of lottery winners. Dependent variables measure logged financing activity, including total funding, VC funding, and Series A funding. All specifications use firm and year fixed effects. Standard errors are clustered at the firm level.

Table 7 deepens our understanding of the VC funding effects by again splitting the sample into small and large firms. The table shows that the negative funding effects for current and past lottery winners are strongly driven by small firms. This differentiated analysis is especially informative given the background of our prior results: Small firms are where the withholding mechanism operates most strongly (Hypothesis 1), where the damage of denying H-1B continuations is lowest due to withholding complementary investments (Hypothesis 3)

and now also where VC funding contracts most (Hypothesis 4). The pattern is consistent across all three margins: behavior, persistence, and funding.

This is the empirical signature of the rule-of-law-as-protection argument from Section 2.1. Large incumbents can potentially shield themselves from discretionary policy uncertainty through political resources, legal capacity, and diversified operational options that resource-constrained startups simply do not have. The result is an asymmetry that shows up across the entire causal chain from firm decisions to innovation persistence to financing outcomes.

Table 7. Lagged H-1B Lottery Winners and Exposure to Denial Rates by Firm Size: Funding Outcomes

	Large Firms			***Small Firms***		
	Total Funding	VC Funding	Series A Funding	Total Funding	VC Funding	Series A Funding
Panel A: Current Lottery Winner						
Lottery Winner ×	−8.407	−6.352	1.095	−29.021***	−16.443***	−6.417***
Fraction Denied	(8.329)	(7.174)	(1.716)	(3.159)	(2.474)	(1.305)
Observations	1431	1431	1431	8453	8453	8453
Panel B: One-Year Lagged Lottery Winner						
L1 Lottery Winner ×	−12.001*	−1.945	3.369**	−22.986***	−14.890***	−4.020***
Fraction Denied	(7.065)	(6.100)	(1.567)	(3.220)	(2.383)	(1.017)
Observations	1248	1248	1248	7362	7362	7362
Panel C: Two-Year Lagged Lottery Winner						
L2 Lottery Winner ×	−15.959**	1.866	0.730	−20.430***	−10.291***	−5.241***
Fraction Denied	(7.814)	(6.634)	(1.723)	(3.341)	(2.295)	(1.207)
Observations	1000	1000	1000	5847	5847	5847
Panel D: Three-Year Lagged Lottery Winner						
L3 Lottery Winner ×	−20.252**	−2.884	3.447	−11.927***	−8.701***	−3.152**
Fraction Denied	(8.113)	(6.585)	(3.058)	(3.775)	(2.665)	(1.441)
Observations	777	777	777	4461	4461	4461
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Notes: This table reports specifications interacting current and lagged H-1B lottery-winner indicators with the aggregate H-1B denial rate, separately for large and small firms. Large firms are defined as firms with more than 500 employees; small firms are defined as firms with 499 or fewer employees. L1, L2, and L3 Lottery Winner indicate whether the firm was a lottery winner one, two, or three years prior, respectively. Fraction Denied is an annual aggregate measure equal to the number of lottery winners denied an H-1B visa divided by the total number of lottery winners. Dependent variables measure logged financing activity, including total funding, VC funding, and Series A funding. All specifications use firm and year fixed effects. Standard errors are clustered at the firm level.

6. Conclusion

We exploit exogenous variation from the H-1B lottery and the rise in discretionary denials during the first Trump administration to study how policy uncertainty affects the strategic factor market for high-skilled global talent. The causal benefits of winning the lottery are sharply attenuated under high uncertainty. This attenuation manifests itself empirically in four specific ways: lower patent application gains from winning the H-1B lottery in times of higher H-1B policy uncertainty; lower persistence of past lottery wins on current innovation; less damage from losing H-1B workers through denials of H-1B continuations; and reduced VC funding for current and past H-1B lottery winners. Across all four margins, effects are concentrated on small startups while large incumbents remain relatively insulated.

All of our findings can be understood from a unified perspective: complementary investment under irreversibility, with real options as the key strategic framework (Chintakananda et al., 2024; Ghemawat, 1991; Pindyck, 1991; Trigeorgis and Reuer, 2017). When the predictability of administrative enforcement falls, the option value of deferring co-specialized investment rises, and firms reduce their commitment to investments that depend on the H-1B worker remaining. This logic explains the firm-level attenuation, the broken innovation chains, the preemptive withholding revealed by reduced denial damage, the VC funding contraction, and the firm-size heterogeneity in a single theoretical framework. The contribution to strategic factor market theory (Barney, 1986; Peteraf, 1993) is to show that the cornerstones of competitive advantage rest on an implicit rule-of-law foundation, which is reshaped when discretionary administrative action erodes the predictability of rules even while the rules themselves remain unchanged.

Our findings have direct implications for the broader question of how democratic institutions support business activity. The rule of law is often discussed as an abstract institutional ideal in debates about democratic resilience. Our results provide a tangible empirical illustration of what its erosion costs: when discretionary administrative action undermines predictable rules, the costs fall disproportionately on the startups that democracies most need to enable future innovation, and on the entrants providing competitive checks on politically connected incumbents. Future research could examine analogous patterns in other strategic factor markets where administrative discretion shapes resource access, and could investigate which firm strategies most effectively buffer against rule-of-law erosion. A central question is whether the political resources that insulate incumbents — lobbying, legal capacity, regulatory engagement — should themselves be considered strategic assets under conditions of elevated executive discretion, with implications for how firms position themselves in democracies whose institutional foundations are increasingly contested.

References

Abel, Andrew B, Janice C Eberly. 1994. American economic association. *The American Economic Review*, **84**(5) 1369–1384.

Adner, Ron, Daniel A Levinthal. 2004. What is not a real option: Considering boundaries for the application of real options to business strategy. *Academy of management review*, **29**(1) 74–85.

Angrist, Joshua D, Jörn-Steffen Pischke. 2009. *Mostly harmless econometrics: An empiricist's companion*.

Arrow, Kenneth J. 1962. Economic welfare and the allocation of resources for invention. Richard R. Nelson, ed., *The Rate and Direction of Inventive Activity: Economic and Social Factors*. 609–626.

Arts, Sam, Bruno Cassiman, Jianan Hou. 2023. Position and differentiation of firms in technology space. *Management Science*, **69**(12) 7253–7265.

Baker, Scott R, Nicholas Bloom, Steven J Davis. 2016. Measuring economic policy uncertainty. *The quarterly journal of economics*, **131**(4) 1593–1636.

Barney, Jay B. 1986. Strategic factor markets: Expectations, luck, and business strategy. *Management Science*, **32**(10) 1231–1241.

Benner, Mary J, Todd Zenger. 2016. The lemons problem in markets for strategy. *Strategy Science*, **1**(2) 71–89.

Bloom, Nicholas. 2009. The impact of uncertainty shocks. *econometrica*, **77**(3) 623–685.

Chen, Jun, Shenje Hshieh, Feng Zhang. 2021a. The role of high-skilled foreign labor in startup performance: Evidence from two natural experiments. *Journal of Financial Economics*, **142**(1) 430–452.

Chen, Jun, Shenje Hshieh, Feng Zhang. 2021b. The role of high-skilled foreign labor in startup performance: Evidence from two natural experiments. *Journal of Financial Economics*, **142**(1) 430–452.

Chintakananda, Asda, David P. McIntyre, Tony W. Tong. 2024. Real options: Connecting with other perspectives and exploring new frontiers. *Academy of Management Perspectives*.

D'Aveni, Richard A., Giovanni Battista Dagnino, Ken G. Smith. 2010. The age of temporary advantage. *Strategic Management Journal*, **31**(13) 1371–1385.

Department of Homeland Security. 2019. Registration Requirement for Petitioners Seeking To File H-1B Petitions on Behalf of Cap-Subject Aliens. 84 Fed. Reg. 888.

Department of Homeland Security. 2025. H-1B Cap-Subject Petitions Selection Process Update. 90 Fed. Reg. 60864.

Dimmock, Stephen G, Jiekun Huang, Scott J Weisbenner. 2022a. Give me your tired, your poor, your high-skilled labor: H-1b lottery outcomes and entrepreneurial success. *Management Science*, **68**(9) 6950–6970.

Dimmock, Stephen G., Jiekun Huang, Scott J. Weisbenner. 2022b. Give me your tired, your poor, your high-skilled labor: H-1b lottery outcomes and entrepreneurial success. *Management Science*, **68**(9) 6950–6970.

Dixit, Avinash K, Robert S Pindyck. 1994. *Investment under uncertainty*.

Donald J. Trump. 2017. Executive Order 13788: Buy American and Hire American. 82 Fed. Reg. 18837.

Fan, Yang, Lubomir P. Litov, Mu-Jeung Yang, Todd R. Zenger. 2026. The technological uniqueness paradox. *Strategic Management Journal*.

Flammer, C., N. Malhotra, C. Loch. 2025a. Call for papers—management science special issue on business and its interrelationship with democratic resilience, geopolitics, and society in a time of change. *Management Science*.

Flammer, Caroline, Aleksandra Kacperczyk. 2019. Corporate social responsibility as a defense against knowledge spillovers: Evidence from the inevitable disclosure doctrine. *Strategic Management Journal*, **40**(8) 1243–1267.

Flammer, Caroline, Neil Malhotra, Christoph H. Loch. 2025b. Call for Papers: Special Issue on Business and Democratic Resilience. Management Science.

Folta, Timothy B. 1998. Governance and uncertainty: The tradeoff between administrative control and commitment. *Strategic Management Journal*, **19**(11) 1007–1028.

Ghemawat, Pankaj. 1991. *Commitment: The Dynamic of Strategy*.

Glennon, Britta. 2024. How do restrictions on high-skilled immigration affect offshoring? evidence from the h-1b program. *Management Science*, **70**(2) 907–930.

Hadani, Michael, Douglas A. Schuler. 2013. In search of El Dorado: The elusive financial returns on corporate political investments. *Strategic Management Journal*, **34**(2) 165–181.

Hall, Robert E., Dale W. Jorgenson. 1967. Tax policy and investment behavior. *American Economic Review*, **57**(3) 391–414.

Hayashi, Fumio. 1982. Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica*, **50**(1) 213–224.

Hillman, Amy J., Gerald D. Keim, Douglas Schuler. 2004. Corporate political activity: A review and research agenda. *Journal of Management*, **30**(6) 837–857.

Huchzermeier, Arnd, Christoph H. Loch. 2001. Project management under risk: Using the real options approach to evaluate flexibility in R&D. *Management Science*, **47**(1) 85–101.

Jorgenson, Dale W. 1963. Capital theory and investment behavior. *American Economic Review*, **53**(2) 247–259.

Kerr, William R., Ramana Nanda, Matthew Rhodes-Kropf. 2014a. Entrepreneurship as experimentation. *Journal of Economic Perspectives*, **28**(3) 25–48.

Kerr, William R, Ramana Nanda, Matthew Rhodes-Kropf. 2014b. Entrepreneurship as experimentation. *Journal of Economic Perspectives*, **28**(3) 25–48.

Kim, Hyunseob, Howard Kung. 2017. The asset redeployability channel: How uncertainty affects corporate investment. *The Review of Financial Studies*, **30**(1) 245–280.

Leiblein, Michael J., John S. Chen, Hart E. Posen. 2017. Resource allocation in strategic factor markets: A realistic real options approach to generating competitive advantage. *Journal of Management*, **43**(8) 2588–2608.

Loch, Christoph H., Michael E. Solt, Elaine M. Bailey. 2008. Diagnosing unforeseeable uncertainty in a new venture. *Journal of Product Innovation Management*, **25**(1) 28–46.

Mallaby, Sebastian. 2022. *The power law: Venture capital and the making of the new future*.

McGrath, Rita Gunther. 1997. A real options logic for initiating technology positioning investments. *Academy of Management Review*, **22**(4) 974–996.

McGrath, Rita Gunther. 2013. *The End of Competitive Advantage: How to Keep Your Strategy Moving as Fast as Your Business*.

Milgrom, Paul, John Roberts. 1995. Complementarities and fit strategy, structure, and organizational change in manufacturing. *Journal of accounting and economics*, **19**(2-3)

179–208.
Nanda, Ramana, Matthew Rhodes-Kropf. 2017. Financing risk and innovation. *Management Science*, **63**(4) 901–918.
National Foundation for American Policy. 2018. H-1B Denial Rates: Analysis of H-1B Data for FY 2018. Tech. rep., NFAP Policy Brief.
National Foundation for American Policy. 2020. H-1B Denial Rates: Analysis of FY 2019 and FY 2020 Data. Tech. rep., NFAP Policy Brief.
Peteraf, Margaret A. 1993. The cornerstones of competitive advantage: A resource-based view. *Strategic Management Journal*, **14**(3) 179–191.
Pich, Michael T., Christoph H. Loch, Arnoud De Meyer. 2002. On uncertainty, ambiguity, and complexity in project management. *Management Science*, **48**(8) 1008–1023.
Pindyck, Robert S. 1991. Irreversibility, uncertainty, and investment. *Journal of Economic Literature*, **29**(3) 1110–1148.
Ries, Eric. 2011. *The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses*.
Romer, Christina D, David H Romer. 2010. The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks. *American economic review*, **100**(3) 763–801.
Sanford, Anthony, Mu-Jeung Yang. 2022. Corporate investment and growth opportunities: The role of r&d-capital complementarity. *Journal of Corporate Finance*, **72** 102130.
Shavell, Steven. 2004. *Foundations of economic analysis of law*.
Teece, David J. 1986. Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy*, **15**(6) 285–305.
Tobin, James. 1969. A general equilibrium approach to monetary theory. *Journal of Money, Credit and Banking*, **1**(1) 15–29.
Trigeorgis, Lenos, Jeffrey J. Reuer. 2017. Real options theory in strategic management. *Strategic Management Journal*, **38**(1) 42–63.
U.S. Citizenship and Immigration Services. 2017. Rescission of Guidance Regarding Deference to Prior Determinations. Policy Memorandum PM-602-0151.
Williamson, O. 1985. *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*.

Appendix A: A Model of Co-Specialized Investment Under H-1B Policy Uncertainty

This appendix formalizes the argument in Section 2.3.3. We model a firm choosing how to deploy an H-1B worker under uncertainty about the worker's continued availability.

Setup

A firm has acquired an H-1B worker through the lottery and must decide how to deploy them. The firm chooses between two strategies:

- **Commit (C):** invest in firm-specific co-specialized assets around the worker, such as specialized roles, long-horizon projects, firm-specific knowledge transfer. This implies an irreversible investment of $K > 0$.
- **Withhold (W):** assign the worker to modular roles that do not require firm-specific co-investment.

After the firm chooses, nature determines whether the worker's continuation is approved (probability $1-p$) or denied (probability p). The denial probability p depends on the policy regime: p_L under stable rule of law and $p_H > p_L$ under elevated discretionary denial rates.

Payoffs

The firm's payoff depends on its action and nature's realization:

	Approve $(1-p)$	Deny (p)
Commit (C)	$V_C - K$	$-K$
Withhold (W)	V_W	0

where $V_C - K > V_W > 0 > -K$ and $K > 0$. The committed strategy yields higher productive value when the worker remains ($V_C - K > V_W$) but inflicts a sunk cost K when the worker is denied. The withholding strategy yields lower productive value V_W when the worker remains but no sunk cost when denied, since the firm has not built any firm-specific assets to lose.

Optimal strategy

The firm commits if and only if the expected payoff from commitment weakly exceeds the expected payoff from withholding:

$$(1-p)V_C - K \geq (1-p)V_W. \tag{6}$$

Rearranging,

$$(1-p)(V_C - V_W) \geq K, \tag{7}$$

which yields the commitment threshold

$$p^* = \frac{V_C - K - V_W}{V_C - V_W}. \tag{8}$$

The firm commits if $p < p^*$ and withholds if $p > p^*$. As the denial probability rises and crosses the threshold, the firm switches discretely from committing to withholding. Under stable rule of law ($p_L < p^*$) the firm commits; under elevated denial rates ($p_H > p^*$) the firm withholds.

Marginal damage of a denial event

Conditional on a denial occurring, the firm's loss is the gap between its payoff under approval and under denial. Under each regime:

$$\Delta_L = V_C - K - (-K) = V_C \quad \text{(under stable rule of law, firm has committed)} \tag{9}$$

$$\Delta_H = V_W - 0 = V_W \quad \text{(under elevated denial rates, firm has withheld)} \tag{10}$$

Since $V_C > V_W$,

$$\Delta_L = V_C > V_W = \Delta_H. \tag{11}$$

The marginal damage of a denial event is larger under stable rule of law than under elevated denial rates. This is the empirical prediction tested in Hypothesis 3.

Firm-size moderator

Firms differ in scale. The productive value the firm can extract from any single worker is bounded by firm size: a small firm has fewer projects to assign the worker to, fewer customers whose needs can be met, and a smaller addressable market. The disruption from a failed co-investment, by contrast, is largely fixed across firm sizes—reorganizing a team, replacing a specialized role, and absorbing project delays consume a similar amount of organizational attention regardless of firm size.

Formally, let firm size be indexed by $S > 0$ and let productive payoffs scale with S while the sunk cost of failed commitment is invariant: $V_C = S \cdot v_C$, $V_W = S \cdot v_W$, and K constant across firms, with $v_C > v_W > 0$. The commitment threshold is then

$$p^*(S) = \frac{S(v_C - v_W) - K}{S(v_C - v_W)}. \tag{12}$$

Differentiating with respect to S,

$$\frac{\partial p^*}{\partial S} = \frac{K}{S^2 \left[v_C - v_W\right]} > 0. \tag{13}$$

The threshold is increasing in firm size: large firms tolerate higher denial probabilities before switching to withholding. Equivalently, small firms switch to withholding at lower denial probabilities than large firms do.

In the limit, as $S \to \infty$, $p^* \to 1$: very large firms commit even when denial is nearly certain, because the productive value of commitment overwhelms the fixed disruption cost. As $S \to 0$, $p^* \to 0$: very small firms withhold even when denial is unlikely, because the productive value of any single worker is small relative to the cost of organizational disruption from a failed investment.

This generates the firm-size prediction in Hypothesis 3: at any given level of policy uncertainty, the withholding response is more likely at small firms, and the gap between Δ_L and Δ_H should therefore be larger for small (startup) firms than for large incumbents.